REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
11-30-2006	Final	06-01-2003 – 11-30-2006

4. TITLE AND SUBTITLE
Pattern-Theoretic Foundations of Automatic Target Recognition in Clutter

5a. CONTRACT NUMBER

5b. GRANT NUMBER
F49620-03-1-0340

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Aaron Lanterman

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Georgia Tech Research Corp.
Georgia Institute of Technology
505 Tenth Street, NW
Atlanta, GA 30332-0420

8. PERFORMING ORGANIZATION REPORT NUMBER
1

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Office of Scientific Research
Dr. Jon Sjogren
875 North Randolph Street
Suite 325 Room 3112
Arlington, VA 22203

10. SPONSOR/MONITOR'S ACRONYM(S)

11. AFRL-SR-AR-TR-07-0186

12. DISTRIBUTION / AVAILABILITY STATEMENT
N/A
Distribution Statement A: unlimited

13. SUPPLEMENTARY NOTES
N/A

14. ABSTRACT
This effort advanced the art of applying Grenander's pattern theory to automatic target recognition (ATR) problems. We extended jump-diffusion ATR algorithms to accommodate unknown infrared camera calibration effects and include more numerically stable diffusion procedures for pose refinement, and developed flexible shape models to accommodate clutter. We also developed performance bounds on estimation and recognition performance for low-frequency radar data, single-image laser radar data, and 3-D ``point cloud'' data assembled from multiple sources. Further work explored data fusion using the ``probability hypothesis density'' approach.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT Unclassified	**b. ABSTRACT** Unclassified	**c. THIS PAGE** Unclassified	None		Aaron Lanterman
					19b. TELEPHONE NUMBER (include area code) 404-385-2548

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

Pattern-Theoretic Foundations of Automatic Target Recognition in Clutter AFOSR grant F49620-03-1-0340 Final Report

Principal Investigator: Aaron Lanterman

School of Electrical and Computer Engineering

Georgia Institute of Technology, Mail Code 0250

Atlanta, GA 30332

(314)-385-2548, lanterma@ece.gatech.edu

May 24, 2007

Abstract

This effort advanced the art of applying Grenander's pattern theory to automatic target recognition (ATR) problems. We extended jump-diffusion ATR algorithms to accommodate unknown infrared camera calibration effects and include more stable diffusion procedures for pose refinement, and developed flexible shape models to accommodate clutter. We also developed performance bounds on estimation and recognition performance for low-frequency radar data, single-image laser radar data, and "point cloud" 3-D data assembled from multiple sources. Further work explored data fusion using the "probability hypothesis density" approach.

Contents

1   Objectives

This research project seeks novel applications of Grenander's pattern theory to problems of Automatic Target Recognition (ATR) in clutter. Model-based algorithms can be powerful, but they can also be fragile if there is a substantial mismatch between the reality and the model. Previous applications of pattern theory to ATR for infrared and laser radar data have considered scenes consisting of targets from a known library against a simple background. While this may be sufficient for relatively simple scenarios, such as tanks against a desert background, such simple "target/background" parameterizations will have difficulty with more cluttered scenes. One fundamental aspect of studying such systems is development of fundamental lower bounds on their performance; such bounds may be used to optimize system parameters.

1.1 Philosophy

Our pattern-theoretic approach to ATR might be thought of as "recognition through simulation" or "recognition through synthesis." In traditional ATR development, simulation plays an important role in the creating numerous data sets to test traditional algorithms. In contrast, in our Grenander-inspired approach, simulation plays a role inside the ATR algorithm itself.

In the field, an ATR system collects some data about an underlying scene of interest. Ideally, we would have a Magic Sensor that says "there is a T62 at this latitude and longitude, and an M60 at this other latitude and longitude." Our real sensors must collect data subject the vagaries of nature and the sensor. In the case of laser radar, we see at targets through the effects of obscuration and perspective projection. The sensor

2

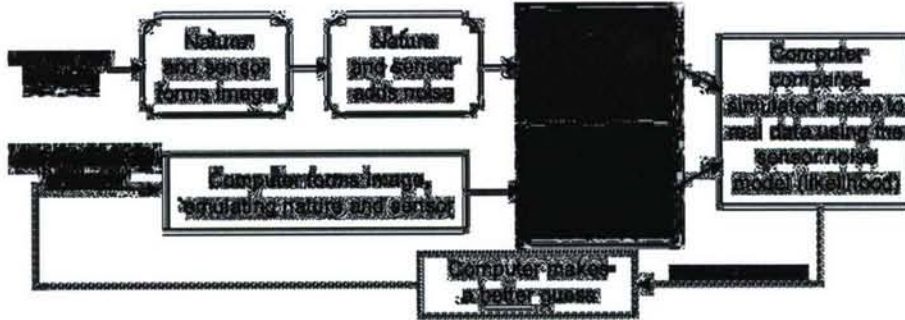adds additional complications such as sampling effects, range noise, and anomalous ("garbage") pixels.

Figure 1: High-level view of our approach to designing ATR algorithms.

As illustrated in Figure 1, our approach to ATR involves simulating a scene that would be generated by a hypothesized set of targets (including types, positions, orientations, etc.), but not including noise. The effect of noise on the real data is encapsulated in the "sensor likelihood" function, which compares our hypothetical computer-generated image with the real data. If the likelihood is high, then the computer-generated image and the real data are a close match, and our hypothesis was a good one. However, our initial hypothesis is probably not that good, and our hypothesized scene does not match the real data well; hence, we invoke a feedback loop, by which the hypothesis is refined, and new hypothesized target sets are formed. The process repeats, with the algorithm continually trying new target types, refining their positions and orientations, etc., trying to get its simulated scene to most closely match the real data.

Note that this process just involves simulating scenes and comparing them to the real data with a sensor likelihood. There are no separate stages of denoising, edge detection, feature extraction, etc. as often found in traditional ATR algorithms.

Although Figure 1 uses laser radar as an example, this philosophy may be used on any sort of sensor, as long as a model for generating data with that sensor is available. Having accurate sensor models, particularly for tasks like scene generation, is important. In this approach, it is vital that we get to know our sensors.

1.2 Context of Effort

Work on this Pattern Theoretic ATR (PTATR) effort under AFOSR funding began in mid-August, 2003. Although the initial effort primarily focused on ATR with infrared data, it split into four related threads. The first two, ATR with infrared data (Section 2.1) and ATR with 3-D data (Section 2.2), constitute the bulk of the PTATR effort. The second two tasks, the creation of performance bounds (Section 2.3) and data fusion using finite-set statistics (Section 2.4), constitute synergistic adjuncts with other efforts. Additional effort was expended towards developing sensor simulation software for use

by the general ATR community, as described in Section 5.2. A subcontract from Jacobs Sverdrup allowed Jason Dixon to spend some time working directly with personnel at Wright-Patterson Air Force Base during the development of the LADAR simulator.

To best leverage the government's support and give AFOSR maximal impact for its investment, work in this area has continued past the September 2006 end date via discretionary funds provided to Aaron Lanterman through the Demetrius T. Paris Junior Professorship.

2 Accomplishments

2.1 ATR with Infrared Data

As described in Appendix F, we are exploring different ways of extending previous models by supplementing the 3-D CAD models of specific targets with flexible models that can accommodate clutter objects not found in the algorithm's target library. Some emphasis has been placed on moving away from strict diffusion implementations, since diffusions involve difficult choices involving step sizes (both in terms of discretizing the diffusion itself and in terms of numerical computations of derivatives); as described in Section 2.1.1, we have explored achieving the effect of diffusions via "little jumps" that avoid some of these complications.

In jump-diffusion algorithms for Bayesian inference, the jumps are typically responsible for handling large changes in the hypothesized configuration, such as the addition or deletion of a target or a change of target type, while the diffusions refine continuous parameters, such as position and orientation. Various kinds of jumps are illustrated in Figure 2.
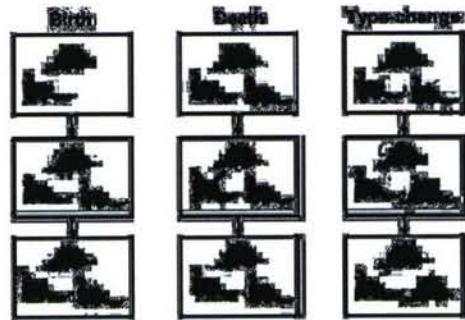

Figure 2: Illustration of jump moves to add hypothesized targets, remove hypothesized targets, and change hypothesized target types.

Our previous work exploring jump-diffusion algorithms for automatic target recognition from infrared scenes assumed that the radiant characteristics of the objects were

4

Figure 3: Flowchart of legacy jump-diffusion algorithm for pattern-theoretic ATR with infrared data.

known. The radiant characteristics of real objects varies widely depending on environmental conditions and operational history; for instance, whether the engine has been running and how hard, whether it is a cloudy or sunny day, whether the gun has just been fired, etc. To accommodate this thermal variability within the Grenander framework, we have incorporated eigentank models into the jump-diffusion code. The eigentanks are defined over the full 3-D surface of the target, although only a portion of that surface is projected onto the detector. The eigentank models are derived from a principal component analysis of thermal profiles simulated under a wide range of conditions, and characterize how different parts of a target "heat up" or "cool down" together. Figure 3 gives a flowchart of our jump-diffusion algorithm for ATR with infrared data, as it roughly stood before the beginning of our PTATR effort. Figure 4 shows where the new advances developed under this effort, including modifications

described in Appendix B and Sections 2.1.1 and 2.1.2, fit into the overall algorithm.

Figure 4: Flowchart illustrating enhancements to the pattern theoretic ATR with infrared data developed under the current effort.

2.1.1 Improved Diffusion Approaches

Background on previous Langevin-based diffusion algorithms: The pattern-theoretic ATR algorithm employed in our studies uses a type of jump-diffusion process[1] to iteratively estimate target characteristics and pose present in a given image. The pose parameters are used to render hypothesized scenes through a set of OpenGL routines, which are compared to the original data image using a likelihood function based on forward-looking infrared (FLIR) or laser radar (LADAR) sensor statistics. The jumps determine the size of the parameter space, i.e., the number of targets, while diffusions update pose parameter estimates.

During the intervals between jumps, the diffusion process takes over and adjusts the continuous configuration pose parameters ($[x, y, \theta]$ – two ground-based coordinate positions and an orientation angle) by small amounts to better align the hypothesized targets with the corresponding detected targets in the data. Diffusions are accomplished using the Langevin stochastic differential equation (SDE):

$$dC_N(\tau) = \nabla_{C_N} H[C_N(\tau)|\mathbf{d}]d\tau + \sqrt{2}dW_N, \tag{1}$$

where W_N is a Wiener process and $H[C_N(\tau)|\mathbf{d}]$ is the logposterior for data $\mathbf{d}$ associated with the configuration parameter vector C_N, which contains the configuration parameters for N targets with fixed classes. x and y are ground-based position coordinates, and θ is the target orientation angle. The time index τ refers to a unit of time within the diffusion interval. Once (1) is discretized, τ can simply be thought of as a discrete time index such that a finite number of diffusions will occur between jumps; that number is an exponential random variable.

The derivative needed in solving the Langevin SDE (1) may be computed with a finite difference approximation:

$$\frac{\partial H(c|\mathbf{d})}{\partial c_p} \approx \frac{H(\dots, c_p + \delta, \dots |\mathbf{d}) - H(\dots, c_p - \delta, \dots |\mathbf{d})}{2\delta}, \tag{2}$$

where c_p is a single parameter of the configuration c, δ is some small deviation of the parameter c_p, and the ellipses indicate the remaining parameters are held fixed.

To summarize this process, the jump-diffusion algorithm starts by estimating an approximate location for a target. If we create a scene by rendering a target at this estimated location, which we will call the hypothesis, we will see that the hypothesized target and corresponding target in the data partially overlap. To refine the initial guess, the diffusion process incrementally adjusts the pose parameters, using the information in the image data. When viewing the hypothesis rendered on top of the data, the overlap between the two improves as the estimated pose parameters converge to their true values. See Figures 5 and 6 for examples of the diffusion process.

Problems with previous diffusion algorithms: When simulating the Langevin SDE, two issues arise: the choice of stepsizes for the derivative computation and the

---

[1] U. Grenander and M.I. Miller, "Representations of Knowledge in Complex Systems," *Journal of the Royal Statistical Society, Series B*, Vol. 56, No. 4, pp. 549–603, 1994. Anyone first approaching Grenander's theory should begin with this paper.
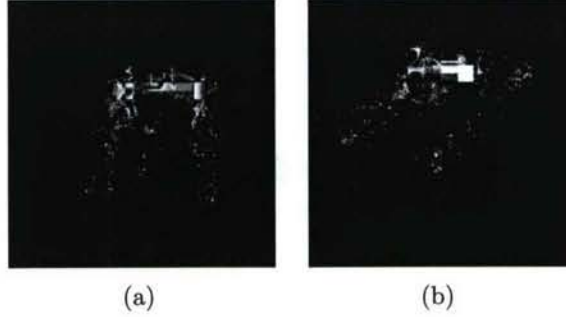
(a)   (b)

Figure 5: Images (a) and (b) contain a sample, synthetic, noisy, LADAR image. A T62 tank sits at the origin of a ground plane. If we assume that our ATR algorithm initially detects a target in the general vicinity of the T62, we may find that our hypothesized T62, denoted by the wire frame outline, does not overlap perfectly with the T62 in the data. This is shown in (a). After a few iterations of the diffusion process, the pose parameters of the hypothesized T62 should match those of the data T62, and the hypothesized T62 should perfectly overlap with the data image, as shown in (b).

(a)   (b)   (c)   (d)
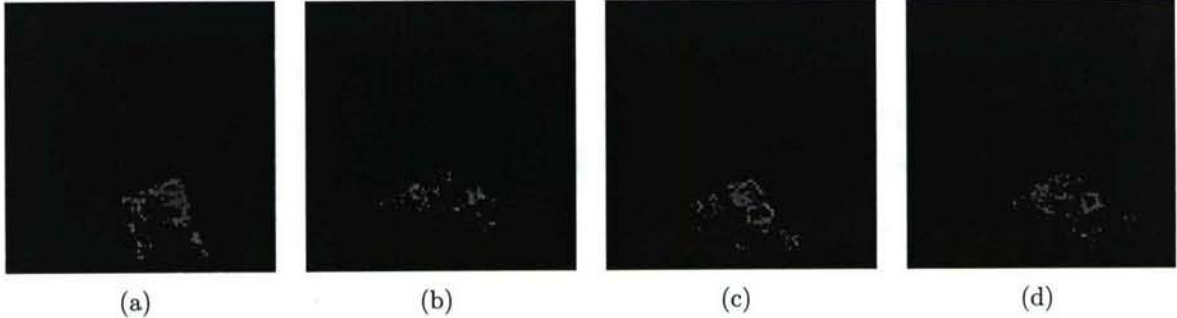
Figure 6: Images (a), (b), and (c) show a hypothesized target rendered over sample FLIR image data. In each of these images, the estimated pose parameters do not match the true values for the corresponding target within the data image, but there is some overlap. After a sufficient number of iterations of the diffusion process, the estimated pose parameters will adjust until there is a closer match with the data, as shown in image (d).

choice of stepsizes for the discretized diffusions themselves. In the jump-diffusion experiments, both of these were determined empirically through a trial-and-error approach that yielded the best adjustments to the configuration parameters. Ideally, these should be automatically determined, but this is problematic because they depend on the types of targets in the scene, the scene's viewing parameters, and the data likelihood function.

Also, Langevin-style diffusions are more natural when there is an analytic solution to the likelihood derivative. As defined in the ATR problem, this derivative must be approximated with the finite difference equation shown in (2). The logposterior H is effectively a function of an image, and we are taking the derivative of this function with respect parameters used to generate the image. The nature of this approximation demonstrates another reason why determining ideal stepsizes for $d\tau$ and dW_N from (1) is not intuitive.

Lastly, diffusions of this form may lead to crude approximations to the pose of detected targets. In some cases, the Langevin diffusion may not converge, but instead oscillate among values close to the true target pose. These characteristics are common to Langevin-style diffusions.[2]

The following discussion will examine our new diffusion algorithm, which redefines the pose parameter refinement problem in way that is more natural to implement and less reliant on arbitrary stepsizes.

New pixel-based diffusion algorithm: In our new diffusion algorithm, we note that the Langevin SDE is not necessary to determine appropriate changes in the coordinate pose parameters (x and y ground plane positions) because those parameters are naturally discretized by the inter-pixel spacing within the image of interest, represented by some real-world unit of measure. For example, if the spacing between neighboring pixels happens to be on the order of centimeters, adjustments to the pose parameters on the order of millimeters will usually not result in a visible change to the hypothesized image. The possibility of subpixel refinement exists, but this requires further study.

The process begins by choosing a *pixel radius*, the desired number of adjacent pixels to consider when determining an updated set of pose parameters for the hypothesized target. A larger pixel radius implies that a larger search space will be considered, increasing the computation time per iteration but decreasing the number of iterations necessary to reach the optimal pose values. Once a pixel radius is chosen, the algorithm proceeds by adjusting the pose parameters for the hypothesized target, selecting values that would result in the image of the hypothesized target moving by a single pixel, or multiple pixels, across the projected surface of the image (see Figure 7). The mapping from pixel coordinates to position coordinates is a function that is built into the OpenGL rendering system we use pattern-theoretic ATR.

In addition to varying the coordinate parameters, we also must consider the orientation parameter θ. This parameter represents the ground-based rotation angle of the target of interest. Rotations do not fit our pixel-based adjustment model as well as changes to the position coordinates, so we are left with choosing a method that will appropriately sample small deviations of θ. This effectively rotates the hypothesized

---

[2]P.J. Green, "A Primer on Markov chain Monte Carlo," *Complex Stochastic Systems*, Eindhoven, pp. 1–62, 2001.

(a)

(-13.67, 13.67)	(-6.83, 13.67)	(0.00, 13.67)	(6.83, 13.67)	(13.67, 13.67)
(-13.67, 6.83)	(-6.83, 6.83)	(0.00, 6.83)	(6.83, 6.83)	(13.67, 6.83)
(-13.67, 0.00)	(-6.83, 0.00)	(0.00, 0.00)	(6.83, 0.00)	(13.67, 0.00)
(-13.67, -6.83)	(-6.83, -6.83)	(0.00, -6.83)	(6.83, -6.83)	(13.67, -6.83)
(-13.67, -13.67)	(-6.83, -13.67)	(0.00, -13.67)	(6.83, -13.67)	(13.67, -13.67)

(b)

Figure 7: In image (a), a tank is located at some ground-based position that we denote (0.00, 0.00). The tank is imaged by a LADAR sensor positioned 100 m away, with a 30° field-of-view angle, pointing directly toward the tank. If a pixel radius of two were chosen for the diffusion, the 24 pixels surrounding the origin pixel would be selected as candidates for the first iteration. For the image as specified in (a), the 24 test pixels and origin pixel will correspond to the grid of (x, y) pose parameters, in centimeters, found in (b). In this case, the space between each pixel is approximately 6.83 cm.

targets by small amounts in both directions. We have found that specifying a *sweep angle* ϕ_s and *sweep angle stepsize* $d\phi_s$ is flexible enough to allow the algorithm to converge to pose values matching those of the corresponding target in the data, under many varying conditions.

Once the candidate coordinate positions have been chosen, the posterior probabilities are computed for each candidate. The best candidate is considered to be the choice that maximizes the logposterior around the chosen samples. The best candidate is accepted with probability $\beta(c_{curr}, c_{next})$, calculated using the Metropolis-Hastings approach:

$$\beta(c_{curr}, c_{next}) = \min \left(\frac{\pi(c_{next})r(c_{next}, c_{curr})}{\pi(c_{curr})r(c_{curr}, c_{next})}, 1 \right). \tag{3}$$

The term c_{next} is the proposed state of the configuration, while c_{orig} is the current state. The functions $r(c_{next}, c_{curr})$ and $r(c_{curr}, c_{next})$ are the transition probabilities.[3]. The function $\pi(c)$ is the probability of being in state c, which in this case is derived from the logposterior H. When selecting candidates during an iteration, candidate position and orientation values can be adjusted by adding a noise term dW_N, just as was done in the Langevin SDE approach, to reduce the possibility of becoming trapped in one of the posterior distribution's local maximums.

A summary of the diffusion algorithm is shown in Table 1.

Table 1: Pixel-based diffusion algorithm.

Initialization
 Determine initial pose parameters $c_{next} \leftarrow [x_0, y_0, \theta_0]$
$\leftarrow [x_0, y_0, \theta_0]$
 Set pixel radius r_p
 Set sweep angle ϕ_s and stepsize $d\phi_s$
Iteration
 Repeat
  (1) Assign $c_{curr} \leftarrow c_{next}$
  (2) Find the set of pixels $\mathcal{R}_p = p_{xy} : \|p_{xy} - p_{x_0 y_0}\| < r_p$
  (3) Find the $N_\phi = \frac{\phi_s}{d\phi_s}$ angular sweep steps in the set $\mathcal{N}_\phi$
  (4) For all pose parameter coordinates in the space $\mathcal{R}_p \times \mathcal{N}_\phi$
   (a) Compute the logposterior probability H
  (5) Assign c_{next} to the $[x, y, \theta]$ that maximizes H over the space $\mathcal{R}_p \times \mathcal{N}_\phi$
  (6) Accept c_{next} with probability $\beta(c_{curr}, c_{next})$
 Until $c_{curr} = c_{next}$

---

[3]See A.D. Lanterman, "Jump-diffusion algorithm for multiple target recognition using laser radar range data," *Optical Engineering*, Vol. 40, No. 8, pp. 1724–1728, 2001, for details.

2.1.2 Calibration

The original algorithms and techniques for pattern-theoretic ATR were evaluated on synthetic image sets. These image sets were created either directly from known thermal intensity values or indirectly by generating thermal intensity values probabilistically from a set of known thermal intensity values. This was appropriate for testing the effectiveness of the algorithm, but it does not reflect how the algorithm will work on real thermal infrared imagery. Therefore, the calibration problem needs to be addressed. Suppose that the relationship between the model and the image is affine and can be summarized as

$$\Lambda_{i,t} = a\lambda_{i,t} + b, \tag{4}$$

where $\Lambda_{i,t}$ is the intensity for target t at intensity region i in the image acquired by the FLIR sensor, a and b are the calibration coefficients, and $\lambda_{i,t} = \sum_j \alpha_{j,t} \Phi_{ij,t} + m_{i,t}$. The technique for estimating expansion coefficients presented in Appendix B can be expanded to include the additional calibration terms. This now creates a nonlinear relationship among the parameters to be investigated, so a new solution must be derived. The analysis follows the previous derivation for the logposterior for pixels on target in terms of the expansion coefficients given in Appendix B, except here we include the calibration coefficients a and b as well. This new logposterior may be written as

$$H(\boldsymbol{\alpha}, a, b | D) = -\sum_t \sum_i \sum_{k \in R_{i,t}} \frac{(\Lambda_{i,t} - d(k))^2}{2\sigma^2} - \sum_t \sum_j \frac{\alpha_{j,t}^2}{2\gamma_{j,t}} \tag{5}$$

$$= -\sum_t \sum_i \sum_{k \in R_{i,t}} \frac{\Lambda_{i,t}^2 - 2\Lambda_{i,t}d(k) + d^2(k)}{2\sigma^2} - \sum_t \sum_j \frac{\alpha_{j,t}^2}{2\gamma_{j,t}} \tag{6}$$

$$= -\sum_t \sum_i \frac{N_{i,t}\Lambda_{i,t}^2 - 2\Lambda_{i,t}\sum_{k \in R_{i,t}} d(k) + \sum_{k \in R_{i,t}} d^2(k)}{2\sigma^2} - \sum_t \sum_j \frac{\alpha_{j,t}^2}{2\gamma_{j,t}} \tag{7}$$

$$= -\sum_t \sum_i \frac{N_{i,t}\Lambda_{i,t}^2 - 2\Lambda_{i,t}D_{i,t} + D_{i,t}}{2\sigma^2} - \sum_t \sum_j \frac{\alpha_{j,t}^2}{2\gamma_{j,t}}. \tag{8}$$

Incorporating $\Lambda_{i,t} = a\lambda_{i,t} + b$ where $\lambda_{i,t} = \sum_j \alpha_{j,t} \Phi_{ij,t} + m_{i,t}$, we must take the derivatives of $H(\boldsymbol{\alpha}, a, b | D)$ with respect to each $\alpha_{j,t}$, a, and b. To make the following derivations cleaner, we begin by mentioning following derivatives:

$$\frac{\partial}{\partial \alpha_{j,t}} \Lambda_{i,t'} = a\frac{\partial}{\partial \alpha_{j,t}} \lambda_{i,t'} = \begin{cases} a\Phi_{ij,t} & \text{if } t = t' \\ 0 & \text{if } t \neq t' \end{cases} \tag{9}$$

$$\frac{\partial}{\partial a} \Lambda_{i,t} = \lambda_{i,t} \tag{10}$$

$$\frac{\partial}{\partial b} \Lambda_{i,t} = 1. \tag{11}$$

We continue with the derivation by computing the derivatives of $H(\boldsymbol{\alpha}, a, b|D)$ with respect to each $\alpha_{j,t}$:

$$\frac{\partial H}{\partial \alpha_{j,t}} = -\sum_{t'}\sum_i \frac{2N_{i,t'}\Lambda_{i,t'}\frac{\partial}{\partial \alpha_{j,t}}\Lambda_{i,t'} - 2\frac{\partial}{\partial \alpha_{j,t}}\Lambda_{i,t'}D_{i,t'}}{2\sigma^2} - \frac{\alpha_{j,t}}{\gamma_{j,t}} \tag{12}$$

$$= -\sum_i \frac{N_{i,t}\left[a\left(\sum_k \alpha_{k,t}\Phi_{ik,t} + m_{i,t}\right) + b\right]a\Phi_{ij,t} - a\Phi_{ij,t}D_{i,t}}{\sigma^2} - \frac{\alpha_{j,t}}{\gamma_{j,t}} \tag{13}$$

$$= -\sum_i \frac{\left(aN_{i,t}\sum_k \alpha_{k,t}\Phi_{ik,t} + aN_{i,t}m_{i,t} + N_{i,t}b\right)a\Phi_{ij,t} - a\Phi_{ij,t}D_{i,t}}{\sigma^2} - \frac{\alpha_{j,t}}{\gamma_{j,t}} \tag{14}$$

$$= -\frac{\sum_i a^2\Phi_{ij,t}N_{i,t}\sum_k \alpha_{k,t}\Phi_{ik,t} + \sum_i \Phi_{ij,t}\left(a^2N_{i,t}m_{i,t} + abN_{i,t} - aD_{i,t}\right)}{\sigma^2} - \frac{\alpha_{j,t}}{\gamma_{j,t}} \tag{15}$$

$$= -\frac{a^2}{\sigma^2}\sum_k \alpha_{k,t}\sum_i N_{i,t}\Phi_{ik,t}\Phi_{ij,t} + \frac{a^2}{\sigma^2}\sum_i \Phi_{ij,t}\left(N_{i,t}m_{i,t} + \frac{bN_{i,t}}{a} - \frac{D_{i,t}}{a}\right) - \frac{\alpha_{j,t}}{\gamma_{j,t}} \tag{16}$$

To maximize the logposterior, we must satisfy these $\sum_t J_t$ necessary conditions:

$$-\frac{a^2}{\sigma^2}\sum_k \alpha_{k,t}\sum_i N_{i,t}\Phi_{ik,t}\Phi_{ij,t} + \frac{a^2}{\sigma^2}\sum_i \Phi_{ij,t}\left(N_{i,t}m_{i,t} + \frac{bN_{i,t}}{a} - \frac{D_{i,t}}{a}\right) - \frac{\alpha_{j,t}}{\gamma_{j,t}} = 0, \ \forall j, t \tag{17}$$

$$\sum_k \alpha_{k,t}\sum_i N_{i,t}\Phi_{ik,t}\Phi_{ij,t} - \sum_i \Phi_{ij,t}\left(N_{i,t}m_{i,t} + \frac{bN_{i,t}}{a} - \frac{D_{i,t}}{a}\right) + \frac{\sigma^2\alpha_{j,t}}{a^2\gamma_{j,t}} = 0, \ \forall j, t \tag{18}$$

Fortunately these are T sets of linear equations, one set for each target, conveniently expressed in matrix form:

$$\left(\begin{bmatrix} \sum_i N_{i,t}\Phi_{i1,t}^2 & \cdots & \sum_i N_{i,t}\Phi_{iJ,t}\Phi_{i1,t} \\ \vdots & & \vdots \\ \sum_i N_{i,t}\Phi_{i1,t}\Phi_{iJ,t} & \cdots & \sum_i N_{i,t}\Phi_{iJ,t}^2 \end{bmatrix} + diag\left(\begin{array}{c} \frac{\sigma^2}{a^2\gamma_{1,t}} \\ \vdots \\ \frac{\sigma^2}{a^2\gamma_{J,t}} \end{array}\right)\right)\begin{bmatrix} \alpha_{1,t} \\ \vdots \\ \alpha_{J,t} \end{bmatrix}$$
$$= \begin{bmatrix} \sum_i \Phi_{i1,t}(\frac{D_{i,t}-bN_{i,t}}{a} - N_{i,t}m_{i,t}) \\ \vdots \\ \sum_i \Phi_{iJ,t}(\frac{D_{i,t}-bN_{i,t}}{a} - N_{i,t}m_{i,t}) \end{bmatrix}. \tag{19}$$

If we keep the $\alpha_{j,t}$ terms constant, we can find the derivatives of $H(\boldsymbol{\alpha}, a, b|D)$ with respect to a and b and maximize the logposterior with respect to these terms. The

13

derivative of $H(\boldsymbol{\alpha}, a, b|D)$ with respect to a is

$$\frac{\partial H}{\partial a} = -\sum_t \sum_i \frac{2N_{i,t}\Lambda_{i,t}\frac{\partial}{\partial a}\Lambda_{i,t} - 2\frac{\partial}{\partial a}\Lambda_{i,t}D_{i,t}}{2\sigma^2} \tag{20}$$

$$= -\sum_t \sum_i \frac{N_{i,t}\Lambda_{i,t}\lambda_{i,t} - \lambda_{i,t}D_{i,t}}{\sigma^2} \tag{21}$$

$$= -\sum_t \sum_i \frac{N_{i,t}\left(a\lambda_{i,t} + b\right)\lambda_{i,t} - \lambda_{i,t}D_{i,t}}{\sigma^2} \tag{22}$$

$$= -\sum_t \sum_i \frac{aN_{i,t}\lambda_{i,t}^2 + bN_{i,t}\lambda_{i,t} - \lambda_{i,t}D_{i,t}}{\sigma^2} \tag{23}$$

$$= -\frac{a}{\sigma^2}\sum_t \sum_i N_{i,t}\lambda_{i,t}^2 - \frac{b}{\sigma^2}\sum_t \sum_i N_{i,t}\lambda_{i,t} + \frac{1}{\sigma^2}\sum_t \sum_i \lambda_{i,t}D_{i,t} \tag{24}$$

and the derivative of $H(\boldsymbol{\alpha}, a, b|D)$ with respect to b is

$$\frac{\partial H}{\partial b} = -\sum_t \sum_i \frac{2N_{i,t}\Lambda_{i,t}\frac{\partial}{\partial b}\Lambda_{i,t} - 2\frac{\partial}{\partial b}\Lambda_{i,t}D_{i,t}}{2\sigma^2} \tag{25}$$

$$= -\sum_t \sum_i \frac{N_{i,t}\Lambda_{i,t} - D_{i,t}}{\sigma^2} \tag{26}$$

$$= -\sum_t \sum_i \frac{N_{i,t}\left(a\lambda_{i,t} + b\right) - D_{i,t}}{\sigma^2} \tag{27}$$

$$= -\sum_t \sum_i \frac{aN_{i,t}\lambda_{i,t} + bN_{i,t} - D_{i,t}}{\sigma^2} \tag{28}$$

$$= -\frac{a}{\sigma^2}\sum_t \sum_i N_{i,t}\lambda_{i,t} + \frac{b}{\sigma^2}\sum_t \sum_i N_{i,t} - \frac{1}{\sigma^2}\sum_t \sum_i D_{i,t}. \tag{29}$$

Two maximize this part of the posterior, the derivatives must satisfy these respective conditions:

$$a\sum_t \sum_i N_{i,t}\lambda_{i,t}^2 + b\sum_t \sum_i N_{i,t}\lambda_{i,t} = \sum_t \sum_i \lambda_{i,t}D_{i,t} \tag{30}$$

$$a\sum_t \sum_i N_{i,t}\lambda_{i,t} + b\sum_t \sum_i N_{i,t} = \sum_t \sum_i D_{i,t}. \tag{31}$$

It is immediately apparent that these conditions can be represented in a form of a matrix to solve for a and b. Also note that these set of equations represent the least squares solution to the problem of determining the affine parameters a and b that best fit the derived thermal intensities $\lambda_{i,t}$ from the data intensities $d(k)$.

We now have two sets of equations: one set that maximizes the logposterior H with respect to the α terms when the a and b terms are held constant and another set that maximizes the same logposterior with respect to the a and b terms when the α terms are held constant. Written together, these sets of equations represent a system of nonlinear equations. Many techniques exist to find solutions to such a system, but

these techniques can be difficult to implement for various reasons. Since this system separates so nicely into two systems of linear equations, it is reasonable to suppose that there may be a way to view the problem in terms of iteratively solving these linear equations. If we can determine a "good" initial value, we may consider the algorithm shown in Table 2 as a guide to iteratively estimate the expansion terms and calibration coefficients.

Table 2: Algorithm to compute the expansion coefficients α_j and thermal calibration coefficients a and b.

Initialization

    Compute $N_{i,t}$ and $D_{i,t}$ $\forall$ i, t
    Assign $\alpha_{j,t}^{(n)} \leftarrow 0$ $\forall$ j, t
    Assign $a^{(n)} \leftarrow 1$
    Assign $b^{(n)} \leftarrow 0$

Iteration

    Repeat
      (1) Assign $a^{(n-1)} \leftarrow a^{(n)}$
      (2) Assign $b^{(n-1)} \leftarrow b^{(n)}$
      (3) Assign $\alpha_{j,t}^{(n-1)} \leftarrow \alpha_{j,t}^{(n)}$ $\forall$ j, t
      (4) Solve for the $\alpha_{j,t}^{(n)}$ coefficients using $a^{(n-1)}$ and $b^{(n-1)}$ $\forall$ j, t
      (5) Compute $\lambda_{i,t} = \sum_j \alpha_{j,t}^{(n)} \Phi_{ij,t} + m_{i,t}$ $\forall$ i, t
      (6) Solve for $a^{(n)}$ and $b^{(n)}$ using $\lambda_{i,t}$, $N_{i,t}$, and $D_{i,t}$ $\forall$ i, t
    Until $\left\| [\boldsymbol{\alpha}, a, b]^{(n)} - [\boldsymbol{\alpha}, a, b]^{(n-1)} \right\| < \epsilon$

2.2 ATR with 3-D Data

The exploitation of 3-D has garnered tremendous attention by DARPA and the Air Force. Different algorithms have been proposed by some of the usual suspects, such as MITRE, Alphatech, and Carnegie Mellon. These algorithms have generally involved extraction of features to achieve real-time performance constraints demanded by DARPA's short-term goals. Our interest in the problem runs deeper. Instead of developing yet another real-time algorithm to compete in a shootout with the myriad real-time algorithms that have already been proposed, we are addressing the more fundamental question: what is the best that we could do with data of a particular quality, independent of whatever particular algorithm is used? As in our infrared work, we bypass the feature extraction stage and conduct inference directly from the full available data, since information may be lost in the feature extraction stage. If a feature-based algorithm manages to achieve our theoretical lower bounds, that implies the features chosen constitute sufficient statistics for the inference problem.

We have developed statistical likelihood models for different kinds of 3-D data. For raw laser radar imagery, we employ models based on the underlying physics of the detector. For assembled "point clouds," whose statistics are a complex interaction of the underlying sensor characteristics and the algorithms used to assemble multiple views, we consider a Poisson point process model chosen for its analytical tractability. Our prime emphasis in this work has been on the creation of performance bounds, as described in Sections 2.3.2 and 2.3.3.

The improved "pixel-based" diffusions we developed for infrared data, described Section 2.1.1, are also applicable to jump-diffusion algorithms for target recognition with laser radar data.[4]

2.3 Performance Bounds via Kullback-Leibler Distances, Chernoff Distances, and Fisher Information

We have developed generic performance bounds, based on Stein's lemma, Chernoff bounds, and Fisher information, which aim to be independent of any particular ATR algorithm. Much of this work follows the lines of theory developed by Anuj Srivastava, Ulf Grenander, and Michael Miller.

2.3.1 Performance Bounds with Low Frequency Radar Data

Inspired by the AFOSR DURIP project titled "Integrated Sensing and Computation for Passive Covert Radar, Signals Intelligence, and Other Applications Driven by Moore's Law," we have demonstrated the power of our information-theoretic performance bounds on the specific application of ATR with passive radar data by comparing our predicted performance measures with empirical performance metrics derived from Monte Carlo runs. Appendix D discusses computing such Chernoff bounds based on a Rician model appropriate for low-frequency radar. Such bounds allow us to answer questions such as "how long must we collect data on an aircraft before we can make an recognition decision with a certain probability of correct decision?"

As described in detail in the next section, the Kullback-Leibler distance provides another metric for detection problems. Appendix E presents an approximation of the Kullback-Leibler distance for Rician models, and Appendix I applies this line of reasoning to low-frequency radar.

2.3.2 Performance Bounds with 3-D Point Cloud Data

At the 2005 AFOSR program review in Raleigh, NC, we presented results on performance analysis using Fisher Information and Kullback-Leibler distances for 3-D data using our Poisson "point cloud" model. Our goal here is to formulate likelihood models on point-cloud data, and not features extracted from those point clouds.

This work was inspired by reading the description of DARPA's E3D program, which sought "efficient techniques for rapidly exploiting 3-D sensor data to precisely locate

---

[4]A. D. Lanterman, "Jump-diffusion Algorithm for Multiple Target Recognition using Laser Radar Range Data," *Optical Engineering*, Vol. 40, No. 8, pp. 1724–1728, 2001.

and recognize targets." The BAA for it contained various demands for different stages of the program, such as "The Target Acquisition and Recognition technology areas will develop techniques to locate and recognize articulating, reconfigurable targets under partial obscuration conditions, with an identification probability of 0.85, a target rejection rate less than 5%, and a processing time of 3 minutes per target or less." This naturally leads to some questions: 1) If such a milestone is not reached, *is that the fault of the algorithm or the sensor?* 2) What performance from a particular sensor is necessary to achieve a certain level of ATR performance, *independent of the question of what algorithm is used?*

Theoretical lower bounds on algorithm performance give algorithm designers a goal to shoot for, and provider criteria for various sensor hardware tradeoffs. In a similar vein, ATR algorithms are typically designed under current computational hardware constraints; however, computers keep getting faster, so it makes sense to ask what ultimate underlying limits are placed by the sensor hardware. The goal is to understand the information content available in the data itself, instead of skipping straight to algorithm development.

Grenander's Pattern Theory provides our philosophical framework. Most 3-D data ATR algorithms (and ATR algorithms with other kinds of data) seek features that are invariant to pose (position and orientation). In contrast, the Grenander approach does not hide from the pose parameter, and explicitly co-estimates it or integrates it out. At a given viewing angle, Target A at one orientation may look much like Target B at a different orientation. As Grenander, Miller, and Srivastava note,[5] "the nuisance parameter of orientation estimation plays a fundamental role in determining the bound on recognition."

A "true" statistical model, which exactly matches the complex interactions of the sensors and Jigsaw-like software, is probably analytically intractable (if it could be developed at all). We base our models on inhomogeneous Poisson processes, since they possess many convenient properties. Although the real data will be distributed more uniformly than a Poisson distribution will predict, the Poisson-based likelihoods should provide useful results. Since modern 3-D point clouds will result from the assembly of multiple views, solely modeling range measurement error across a single line of sight is be insufficient. We assume that the points are seen with an additive Gaussian error of appropriate covariance. This corresponds to a "translated Poisson process," where we essentially assume that the intensity of the observed Poisson process is given by visible portion of the shell of the model convolved with a "point spread function" defined by the measurement covariance.

The experiments described in this section employed four target models, shown in Figure 8, taken from the AFRL 3D Challenge Problem, which was distributed on DVDs at the 2003 SPIE Defense & Security Symposium. The performance graphs in the remaining subsections plot various metrics vs. the standard deviation of the measurement errors. A circularly symmetric error density is assumed.

The DVDs for the 2003 Challenge Problem only have five different look angles per

---

[5]U. Grenander, M.I. Miller, and A. Srivastava, "Hilbert-Schmidt Lower Bounds for Estimators on Matrix Lie Groups for ATR," *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 20, No. 2, Aug. 1998, pp. 790–802.
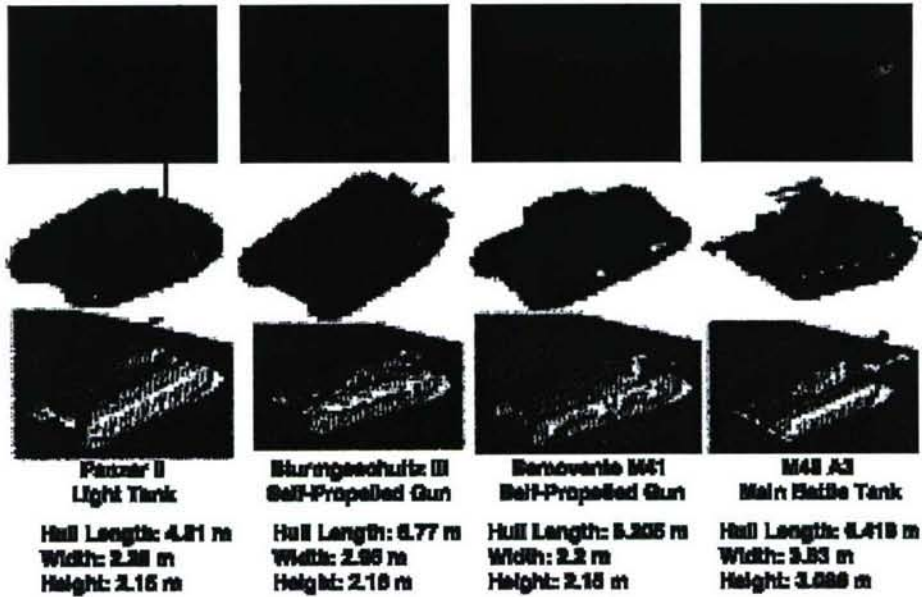
Figure 8: Four targets from the AFRL 2003 3D Challenge Problem, used in our 3-D point-cloud performance model experiments.

target. Hence, we computed bounds for one particular look angle, using adjacent look angles to compute derivatives with respect to orientation. The computed information metrics will, in general, be a function of the orientation, since targets will look different to the sensor at different orientations (in particular, some parts of the target may be obscured if no views are available from a certain range of angles.) The limitations of this Challenge Problem data set was one of the motivating factors of the development of our new simulator described in Appendix A.

Cramér-Rao bounds for point-cloud data: Cramér-Rao bounds for continuous pose parameters are given by the diagonal elements of the inverse of the Fisher information matrix. Cross-terms show how estimate errors are correlated. Figures 9, 10, and 11 illustrate Cramér-Rao bounds on the pose parameters of x-position, y-position, and orientation angle, respectively, for the Sturmgeschultz and Semovente targets. We assume that the target is sitting on a flat surface, and hence the z-position is known. For each target type, two lines are given. One line corresponds to performance in an artificial case where the all of the parameters are known except the target of interest. The other shows performance in the case where all three parameters need to be co-estimated. The distance between the lines shows the "performance hit" that results from the coupling of the parameters.

We emphasize that even if a rather accurate model is available, we would not expect a real system to achieve the bound in practice, since there will always be effects in the real data not present in the model. In addition, in this particular case of 3-D point cloud data, our inhomogenous Poisson model was formulated for computational convenience, not fidelity to reality. We hope, however, that the bounds can provide insight into
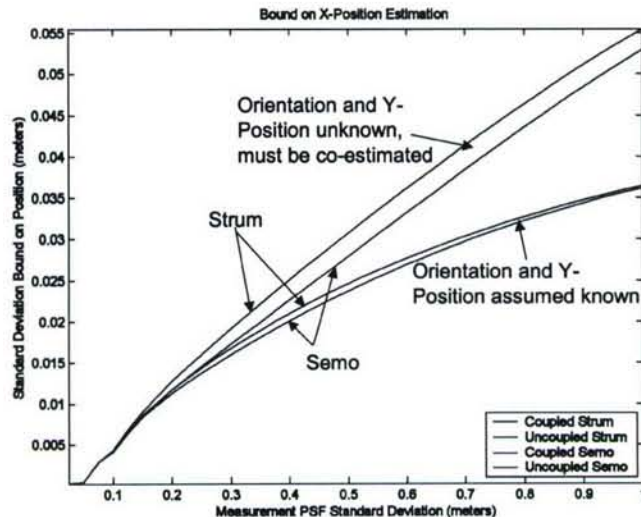
18

Figure 9: Cramér-Rao bounds on the x-position pose parameter of a Sturmgeschulz and Semovente as a function of measurement error.

overall performance trends. In practice, it may be feasible to vary the overall intensity of the Poisson process to bring the bounds into accordance with results from Monte Carlo runs.

I-divergence metrics for point-cloud data: Likelihood ratios are the key statistics in hypothesis testing problems such as automatic target recognition. The Kullback-Leibler distance, also known as the relative entropy, is the expected value of the likelihood ratio, assuming the data is generated according to the "alternative hypothesis" (in the "alternative" vs. "null" nomenclature). According to Stein's lemma, the Kullback-Leibler distance drives the asymptotic performance of the detection problem. For Gaussian data, the Kullback-Leibler distance reduces to a squared-error metric. For Poisson data, as assumed in our point-cloud data model, the Kullback-Leibler distance reduces to Csiszár's I-divergence.

When computing the I-divergence to compare targets, we adjust the pose of the "alternate" target to get closest match to the "true" target as seen by the sensor system. The work of Grenander, Srivastava, and Miller[6] shows that although hypothesis testing with nuisance parameters (in this case) will be dominated by this I-divergence term, there will also be a second term involving the Cramér-Rao bound on the nuisance parameters. This notion links estimation and recognition performance, and connects the experiments described this subsection with the those described in the previous subsection. Here, we only consider the effect of the "first term," i.e. the I-divergence. We hope to study the inclusion of the second term in future work.

Figures 12, 13, 14, and 15 plot the I-divergence metrics (lower I-divergence indicates greater difficulty of discrimination) between the targets. Each of the four plots takes a

---

[6]U. Grenander, A. Srivastava, and M.I. Miller, "Asymptotic Performance Analysis of Bayesian Target Recognition," *IEEE Trans. Information Theory*, Vol. 46, No. 4, July 2000, pp. 1658–1665.
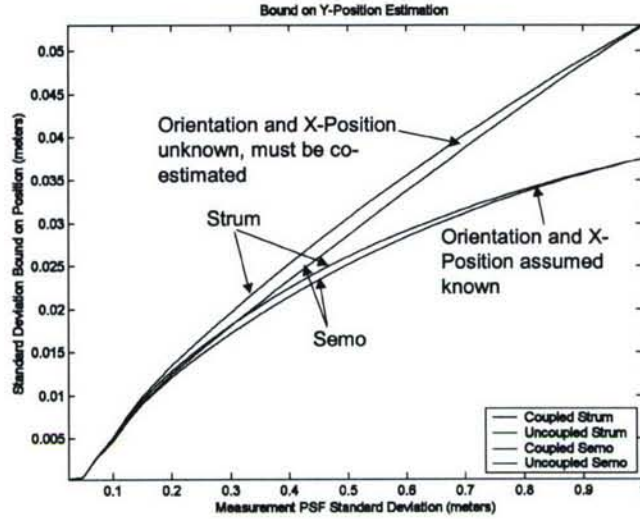
Figure 10: Cramér-Rao bounds on the y-position pose parameter of a Sturmgeschulz and Semovente as a function of measurement error.

single target type to be the "alternative" hypothesis; each line within a plot represents a different "null" hypothesis.

Interestingly, the I-divergence is *not symmetric* in general; for instance, the $I(M48\|panzer)$ line in Figure 15 is different than the $I(panzer\|M48)$ line in Figure 12. This is why we were sure to make a distinction between the "alternative" and "null" hypothesis in the preceeding paragraphs. During the question & answer discussion period after Aaron Lanterman's closing talk at the *ARO/ARMDEC Workshop on Information-Theoretic Image Processing*, Prof. Al Hero of the Univ. of Michigan commented that we should not shy away from this asymmetry or find it unusual; instead, we should "embrace the asymmetry."

Such asymmetry manifests itself in the "pop-out" experiments in psychology.[7] Consider setting a probability of a Type-1 error (usually called "false alarm") in a Neyman-Pearson detection framework. Consider two hypothetical problems, illustrated in Figure 16. In the problem on the left, we want to detect a Panzer in a sea of M48 clutter; in the problem on the right, we want to detect an M48 in a sea of Panzer clutter. It may be counterintuitive, but one problem will be more difficult than the other. Both I-divergence curves for these two target types are shown in Figure 17. Note that the problem illustrated on the left of Figure 16 is easier than that on the right for measurement noise with standard deviation greater than 0.35 meters, whereas the problem illustrated on the right of 16 is easier than that on the left for measurement noise with standard deviation less than 0.35 meters.

---

[7]A.L. Yuille and J.M. Coughlan, "Fundamental Limits of Bayesian Inference: Order Parameters and Phase Transitions for Road Tracking," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 2, No. 2, Feb. 2000, pp. 160–173; Y.N. Wu, S.C. Zhu, and Z. Liu, "Equivalence of Julesz Ensembles and FRAME Models," Intl J. of Computer Vision, Vol. 38, No. 3, 2000, pp. 247–265; see Section 7 in particular.
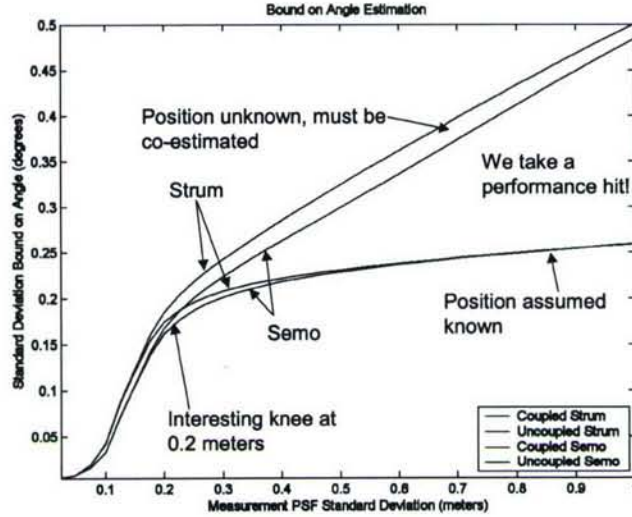
Figure 11: Cramér-Rao bounds on the angle pose parameter of a Sturmgeschulz and Semovente as a function of measurement error.

The cases of Panzer vs. Semovente (Figure 18), and Panzer vs. Sturmgeschultz (Figure 19) have similar crossover points in the 0.3 to 0.4 meter region. The M48 vs. Semovente (Figure 20) and M48 vs. Strumgeschultz cases have lower crossover points, at around 0.2 and 0.1 meters, respectively.

The Semovente vs. Sturmgeschultz case (Figure 22) is particularly interesting, since the two curves lie almost on top of one another; essentially, at any resolution, finding a Semovente in a sea of cluttering Sturmgeschulzen is no more or less difficult than finding a Strumgeschulz in a sea of cluttering Semoventes.

2.3.3   Performance Bounds with Single-Image Laser Radar Data

If we are not trying to use point-cloud data, as in the previous subsection, and are instead single views from a laser radar, more accurate likelihood models may be employed. Appendix C presents results on the computation of Cramér-Rao bounds on pose parameters using such models.

2.4   Finite-Set Statistics for Multitarget Tracking

Inspired by Keith Kastella's work with Joint Multitarget Densities for AFRL, we integrated another research effort (previously supported by startup funds from the School of Electrical Engineering, and primarily supported by funds from the Paris Professorship) on data fusion in multitarget tracking using Ronald Mahler's Finite-Set Statistics, particularly his Probability Hypothesis Densities (PHDs), into this PTATR effort.

Via simulations, we have illustrated the promise of PHD-based multitarget, multisensor tracking using an FM-radio-based passive radar scenario designed to match a system being developed by NATO NC3A. Appendix G addresses some issues that
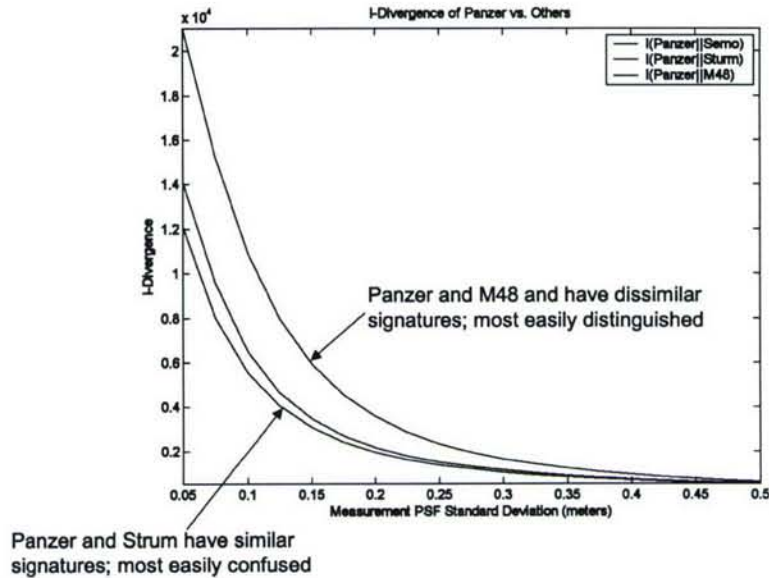
Figure 12: I-divergences between the Panzer model (taken to be the "alternative" hypothesis) and the Semovente, Sturmgeschulz, and Panzer models (taken to be the "null" hypothesis).

have vexed the PHD community, particularly the extraction of peaks (which represent targets) from the PHD, and Appendix H studies the effect of multipath on the algorithm. Our eventual goal is to apply such algorithms to data collected using equipment purchased under the AFOSR DURIP project titled "Integrated Sensing and Computation for Passive Covert Radar, Signals Intelligence, and Other Applications Driven by Moore's Law."

Although our particular example of data fusion and tracking using PHDs happens to be passive radar, the theory is quite general and may be applied to data fusion with other kinds of sensors.

3   Personnel Supported

Jason Dixon, graduate student, fully supported by the AFOSR PTATR grant from August 15, 2004 to August 31, 2006 (41,344.70); support to continue the work past August 2006 provided by the Demetrius T. Paris Professorship. Primary developer of ATR theory and algorithms for target recognition with FLIR and LADAR data, as well as associated applied performance analysis techniques

Lisa Ehrman, graduate student, partially supported by the AFOSR PTATR grant from January 1, 2004 to April 30, 2004 (4,791.08) to focus on performance estimation via Kullback-Leibler distances and Chernoff information measures, with particular application to ATR with low frequency radar. Primarily supported by NATO Consultation, Command and Control Agency (NC3A).
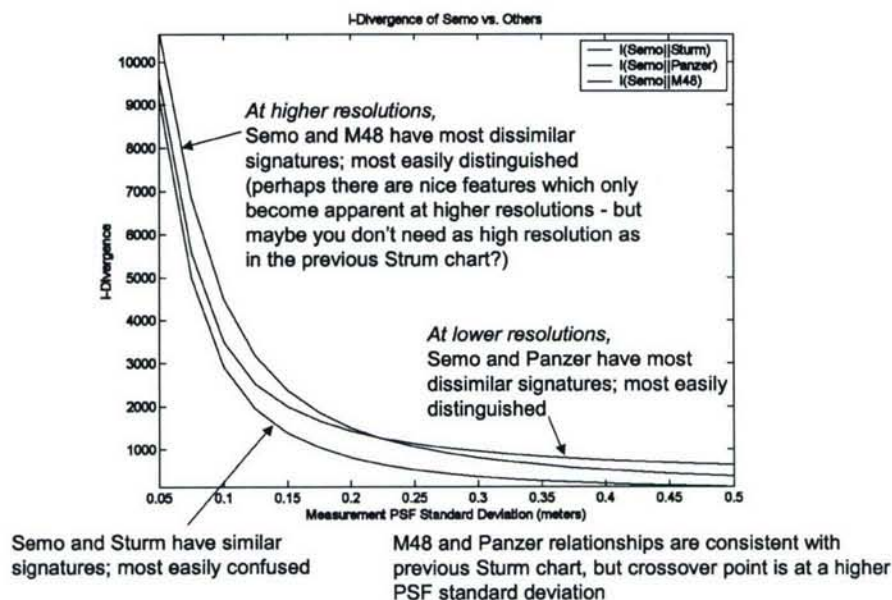
22

I-Divergence of Semo vs. Others

At higher resolutions, Semo and M48 have most dissimilar signatures; most easily distinguished (perhaps there are nice features which only become apparent at higher resolutions - but maybe you don't need as high resolution as in the previous Strum chart?)

At lower resolutions, Semo and Panzer have most dissimilar signatures; most easily distinguished

Semo and Sturm have similar signatures; most easily confused

M48 and Panzer relationships are consistent with previous Sturm chart, but crossover point is at a higher PSF standard deviation

Figure 13: I-divergences between the Semovente model (taken to be the "alternative" hypothesis) and the Sturmgeschulz, Panzer, and M48 models (taken to be the "null" hypothesis).

4 Technical Publications

Some of the publications described below will be available from the PTATR (Pattern-Theoretic Automatic Target Recognition) website at `users.ece.gatech.edu/~lanterma/ptatr`.
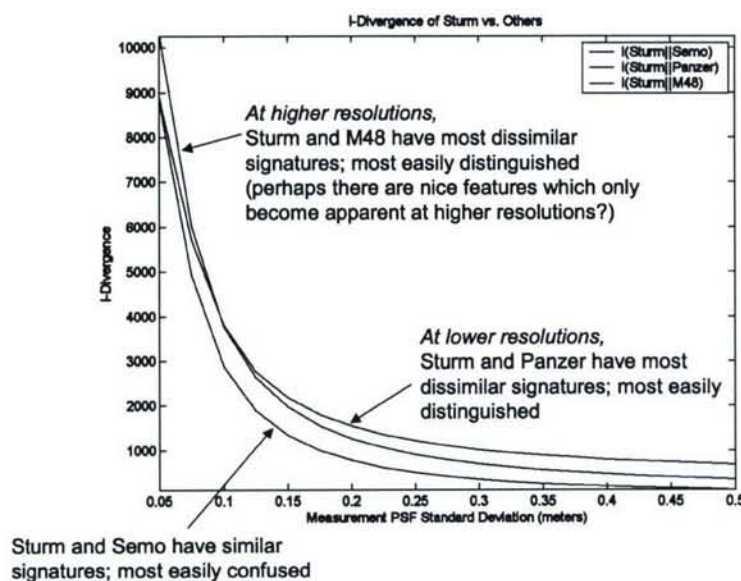
I-Divergence of Sturm vs. Others

At higher resolutions,
Sturm and M48 have most dissimilar
signatures; most easily distinguished
(perhaps there are nice features which only
become apparent at higher resolutions?)

At lower resolutions,
Sturm and Panzer have most
dissimilar signatures; most easily
distinguished

Sturm and Semo have similar
signatures; most easily confused

Figure 14: I-divergences between the Sturmgeschulz model (taken to be the "alternative" hypothesis) and the Semovente, Panzer, and M48 models (taken to be the "null" hypothesis).

A select subset, particularly preprints of submitted or to-be-submitted journal papers, which are referred to in other sections of this document, are provided as appendices to this report. Note that later versions of the papers that may appear on the PTATR website or in print may differ than the early drafts here.

4.1 Doctoral Dissertations

Georgia Tech dissertations may be easily obtained from the Electronic Thesis and Dissertation Collection at `etd.gatech.edu`.

J.H. Dixon, *Pattern-Theoretic Automatic Target Recognition for Infrared and Laser Radar Data*, expected to be completed Summer 2007.

L.M. Ehrman, *An Algorithm for Automatic Target Recognition Using Passive Radar and an EKF for Estimating Aircraft Orientation*, Fall 2005.

M. Tobias, *Probability Hypothesis Densities for Multitarget, Multisensor Tracking with Application to Passive Radar*, Spring 2006.

4.2 Journal Publications

L.M. Ehrman and A.D. Lanterman, "Chernoff-Based Prediction of ATR Performance from Rician Radar Data, with Application to Passive Radar," *Optical Engineering*, letter of acceptance subject to minor revision received Feb. 2, 2006; revision submitted Feb. 2007. (Appendix D).
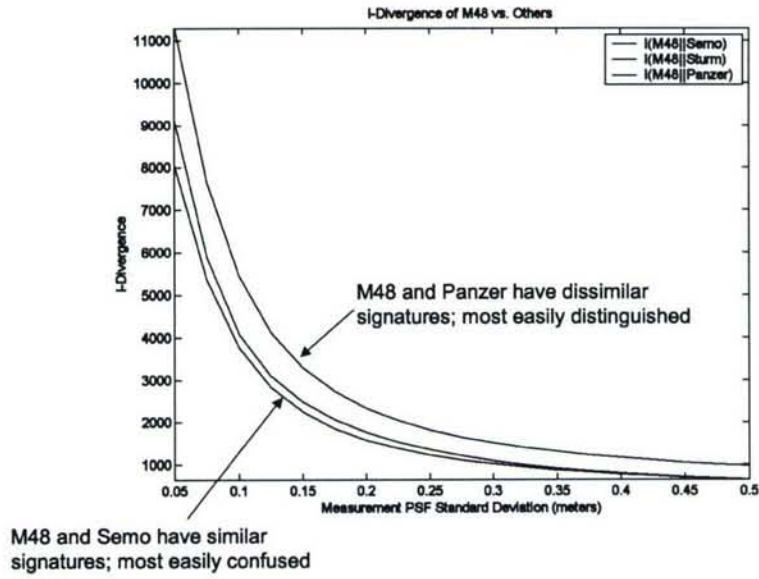
Figure 15: I-divergences between the M48 model (taken to be the "alternative" hypothesis) and the Semovente, Sturmgeschulz, and Panzer models (taken to be the "null" hypothesis).

L.M. Ehrman and A.D. Lanterman, "A Laplace Approximation of the Kullback-Leibler Distance Between Ricean Distributions," *IEEE Trans. on Information Theory*, to be submitted. (Appendix E).

A.D. Lanterman, "Continuous-Time Jump Processes, with Application to Shapes on the Lattice," *Statistics and Computing*, letter requesting major revision received Sept. 2005; manuscript undergoing revision. (Appendix F)

M. Tobias and A.D. Lanterman, "Probability Hypothesis Density-Based Multitarget Tracking with Bistatic Range and Doppler Observations," *IEE Proc. Radar, Sonar, and Navigation*, Vol. 152, No. 3, June 2005, pp. 195–205. (Available from `ieeexplore.ieee.org`, paper number 01459156.)

M. Tobias and A.D. Lanterman, "Techniques for Birth Particle Placement in the PHD Particle Filter, Applied to Passive Radar," *IEE Proc. Radar, Sonar, and Navigation*, submitted April 7, 2007. (Appendix G).

M. Tobias and A.D. Lanterman, "Multipath Effects and the PHD Particle Filter," *IEE Proc. Radar, Sonar, and Navigation*, to be submitted. (Appendix H).

4.3 Conference Publication

L.M. Ehrman and A.D. Lanterman, "Robust Algorithm for Automated Target Recognition using Precomputed Radar Cross Sections, *Automatic Target Recognition XIV*, Proc. SPIE 5426, Ed: F.A. Sadjadi, April 12-16, 2004, pp. 197–208. (Appendix I).
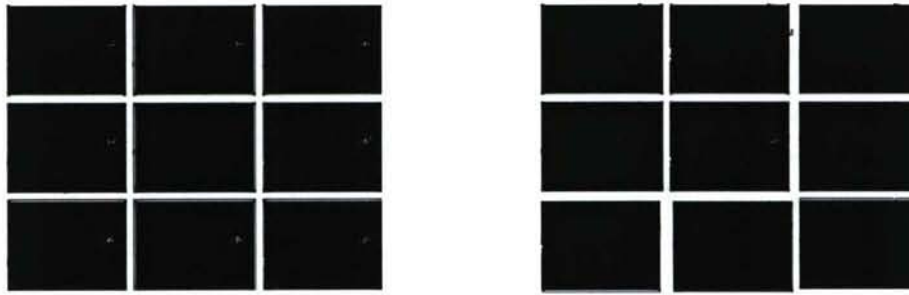
25

Figure 16: Illustration of two asymmetric "pop out experiments." On the left, we want to detect Panzer among cluttering M48s; on the right, we want to detect an M48 among cluttering Panzers.

A.D. Lanterman, "Passive Radar Imaging and Target Recognition using Illuminators of Opportunity," *NATO Symposium on Target Identification and Recognition Using RF Systems*, Oslo, Norway, Oct. 11-13, 2004.

L.M. Ehrman and A.D. Lanterman, "Assessing the Performance of a Covert Automatic Target Recognition Algorithm," *Automatic Target Recognition XV*, Proc. SPIE 5807, Ed: F.A. Sadjadi, Orlando, FL, March 28-April 1, 2005, pp. 77–87.

J.H. Dixon and A.D. Lanterman, "Toward Practical Pattern-Theoretic ATR Algorithms for Infrared Imagery," *Automatic Target Recognition XVI*, SPIE Vol. 6234, Ed: F.A. Sadjadi, April 2006, pp. 212–220. (Appendix B).

J.H. Dixon and A.D. Lanterman, "Information-Theoretic Bounds on Target Recognition Performance from Laser Radar Data," *Automatic Target Recognition XVI*, SPIE Vol. 6234, Ed: F.A. Sadjadi, April 2006, pp. 394–403. (Appendix C).

5 Interactions/Transitions

5.1 Conference and Workshops

J.H. Dixon, "Toward Practical Pattern-Theoretic ATR Algorithms for Infrared Imagery" and "Information-Theoretic Bounds on Target Recognition Performance from Laser Radar Data" (papers listed above under "Conference Publications"), SPIE Defense and Security Symposium, Orlando, FL, April 17-21, 2006 (1,184.04). A.D. Lanterman also attended (1,003.47).

A.D. Lanterman, "Passive Radar Imaging and Target Recognition using Illuminators of Opportunity," *NATO Symposium on Target Identification and Recognition Using RF Systems*, Oslo, Norway, Oct. 11-13, 2004 (trip paid for by NATO).

A.D. Lanterman, "General Pattern Theory," *AFRL ATR Theory MURI Workshop*, Dayton, OH, Dec. 1, 2004 (582.67).
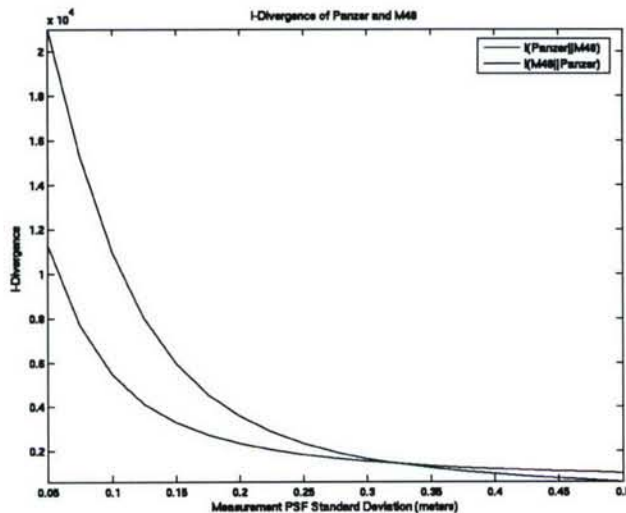
Figure 17: I-divergences between Panzer and M48 models.

A.D. Lanterman, "A Mathematical Theory of Automatic Target Recognition," *Workshop on Quantum Algorithms for Signal, Image, and Data Processing*, UCSD, La Jolla, CA, Dec. 7-9, 2004. Invited talk (trip paid for by UCSD).

A.D. Lanterman (with L.M. Ehrman), "Assessing the Performance of a Covert Automatic Target Recognition Algorithm," SPIE Defense and Security Symposium, Orlando, FL, March 28-April 1, 2005. (See paper reference above).

A.D. Lanterman (with J.H. Dixon and L. Fomundam), "Information-Theoretic Bounds on ATR Performance from Laser Radar Data," *AFOSR 2005 Program Review for Sensing, Imaging and Object Recognition*, NCSU, Raleigh, NC, May 25-27 (681.71). J.H. Dixon also attended (240.00)

A.D. Lanterman, "General Pattern Theory Applied to ATR," *ARO/AMRDEC Workshop on Information-Theoretic Image Processing*, Redstone Arsenal, Huntsville, AL, June 14-15, 2005. Invited talk (trip paid for by Army organizations).

A.D. Lanterman, "A Passive Radar Testbed at Georgia Tech," *4th Multinational Passive Covert Radar Conference*, Syracuse Univ. Hotel and Conference Center, Syracuse, NY, Oct. 5-7, 2005 (797.21).

M. Tobias (with A.D. Lanterman), "Using the Probability Hypothesis Density for Multitarget Tracking with Passive Radar," *4th Multinational Passive Covert Radar Conference*, Syracuse Univ. Hotel and Conference Center, Syracuse, NY, Oct. 5-7, 2005 (653.16).

5.2 Transition: LADAR Simulator

Greg Arnold of AFRL/SNAT became interested in our use of OpenGL in our custom-made laser radar and infrared simulation tools. Some of their previous efforts to cre-
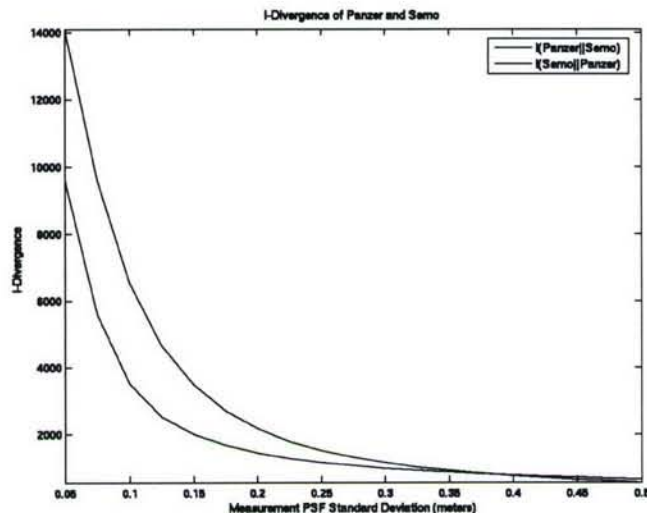
Figure 18: I-divergences between Panzer and Semovente models.

ate synthetic datasets for use by the ATR community involved POV-Ray and IRMA, each of which present various difficulties. We provided him with a documented copy of our simulation code, which he provided to Jeremy Olson and Kyle Erickson of AFRL/SNAA, who performed a detailed review and provided an extensive set of comments, which Jason Dixon used to prepare an updated version of the simulator. We went through several such review/update iterations.

The simulator was originally created to create scenes based on a "heterodyning" laser radars, such as those studied by Jeff Shapiro of MIT. Such radars have a characteristic ambiguity in range manifest in the banded structure of the background. We revised the code, particularly the noise generation code, to allow modeling of the "direct detection" radars studied by T.J. Klausutis of AFRL/MNGI and his colleagues.

Our resulting "LADAR Simulator" code, its associated user's manual, and cross-linked help pages for a small Application Programming Interface (API) that will let users develop their own MATLAB scripts using the simulator's scene generating capabilities is currently available via the web at `users.ece.gatech.edu/\~jdixon/ladar\_sim`. The user's manual is also attached to this report as Appendix A.

In an e-mail dated April 30, 2007, Greg Arnold informed us that his group completed a new Challenge Problem description and data set created using our LADAR Simulator, and that he was reviewing it and working on the "challenge experiments" to send for evaluation for public release.
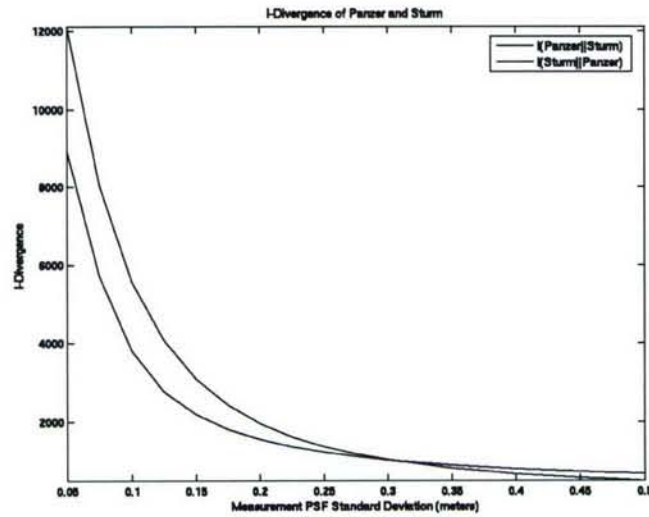
6 Patent Disclosures

None.

Figure 19: I-divergences between Panzer and Sturmgeschultz models.

7 Honors

A.D. Lanterman, Richard M. Bass/Eta Kappa Nu ECE Oustanding Junior Teacher Award (2006), as voted on by the senior class.

A.D. Lanterman, named the Demetrius T. Paris Junior Professor beginning in September 2004. This special three-year Chair position was founded to support the development of young faculty.

A.D. Lanterman, NIC Certificate of Excellence "for outstanding contributions to the National Intelligence Council and exceptional service to the Intelligence Community," April 2001.
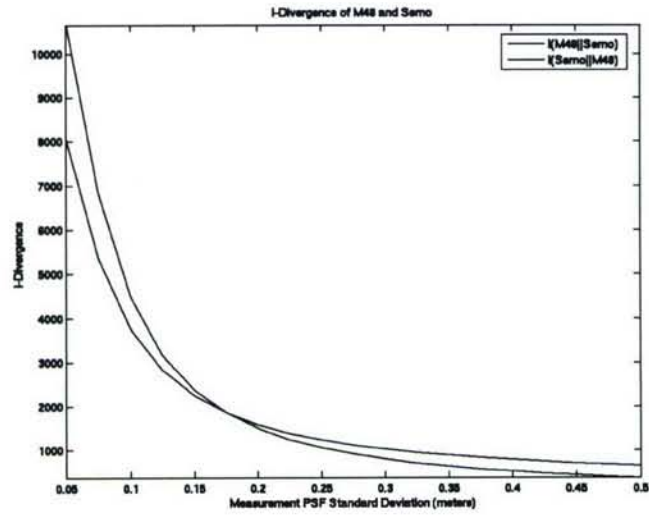
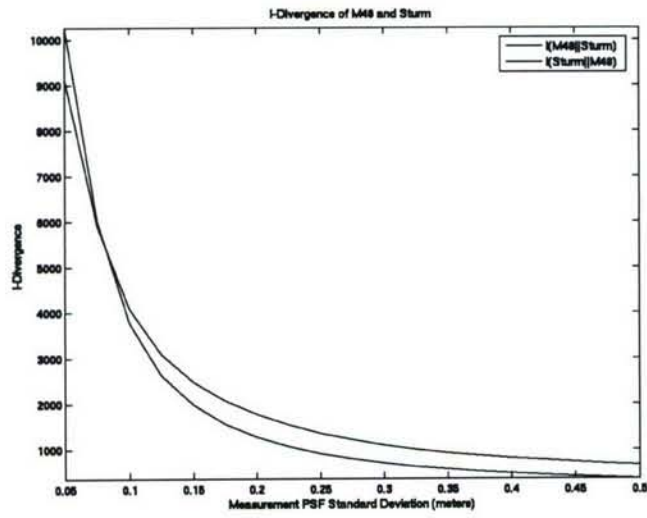Figure 20: I-divergences between M48 and Semovente models.

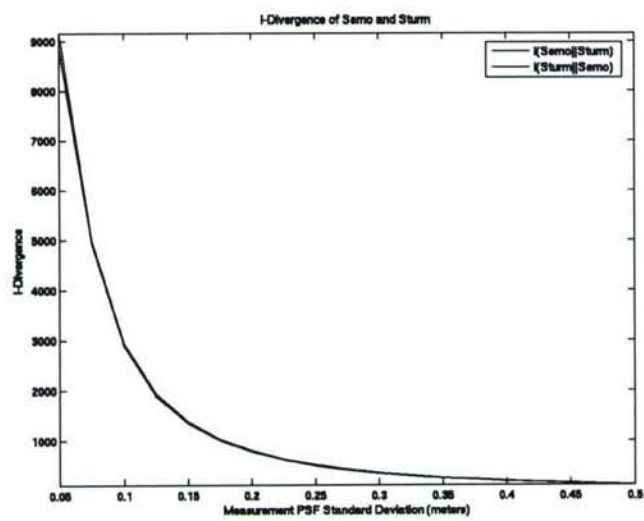Figure 21: I-divergences between M48 and Sturmgeschultz models.

Figure 22: I-divergences between Semovente and Sturmgeschultz models.